White Paper

STREAMLINING VM BACKUPS WITH THE PUBLIC CLOUD

Managing Your VMware, VMC, Hyper-V, and Nutanix Backups With the Public Cloud

Q418-CON-10819

Identifying the Core Issues

The rise of virtualization as a business tool has dramatically enhanced server and primary storage utilization. By allowing multiple operating systems and applications to run on a single physical server, organizations can significantly lower their hardware costs and take advantage of efficiency and agility improvements as more and more tasks become automated. This also alleviates the pain of fragmented IT ecosystems and incompatible data silos.

Currently, this virtualization juggernaut shows no sign of slowing. As businesses recognize the potential for increased reliability and scalability offered by virtual technology, they are ramping up their investments in data center modernization and upgrading. In fact, **33 percent of the respondents** to a recent ESG survey on cloud usage said that making greater use of server virtualization was one of their top five spending priorities for the next 12 to 18 months.

2018 Data Center Modernization Spending Priorities

In which of the following areas of **data center modernization** will your organization make the **most** significant investments over the next 12-18 months? (Percent of respondents, N=544, five responses accepted)

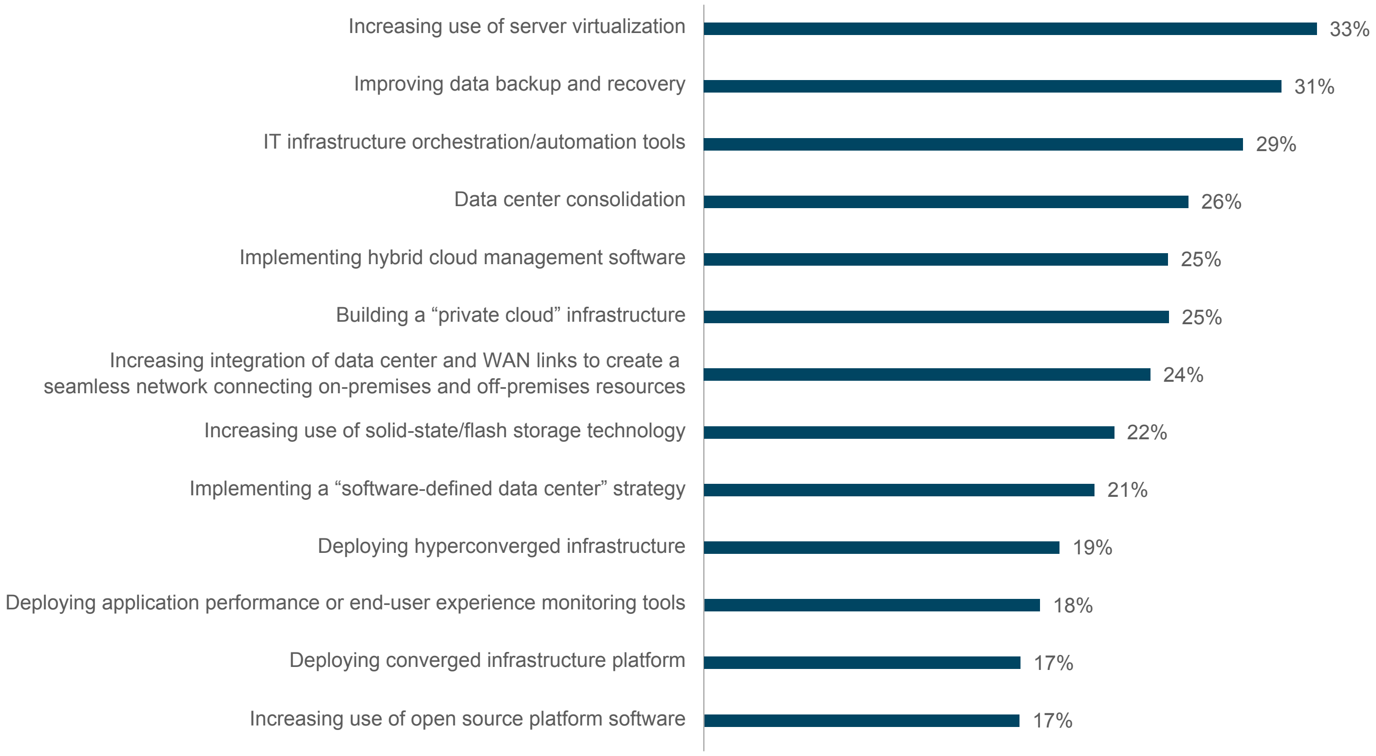

Source: ESG, 2018 IT Spending Intentions Survey

Protecting these virtualized environments and the ever-growing amount of structured and unstructured data being created typically requires a complex, on-premises (on-prem) secondary storage model that imposes heavy administrative overhead and infrastructure costs. However, the increasing pressure on IT teams to maintain business continuity and information governance is changing how businesses view infrastructure resiliency and long-term data retention. As a result, IT teams are seeking new solutions to ensure immediate availability and complete protection of the data that resides within their VMware, VMware Cloud on AWS, Hyper-V, and Nutanix environments.

From Our Customer

In the past, “a typical day at the Port of New Orleans was spending a lot of time with aging infrastructure.” said David Cordell, Chief Technology Officer of the port. “For our backup infrastructure, it would take hours, in some cases days, to retrieve files. We would have to move from one system to another system before we could just restore it to its original location.”

Legacy VM Data Protection Complexity

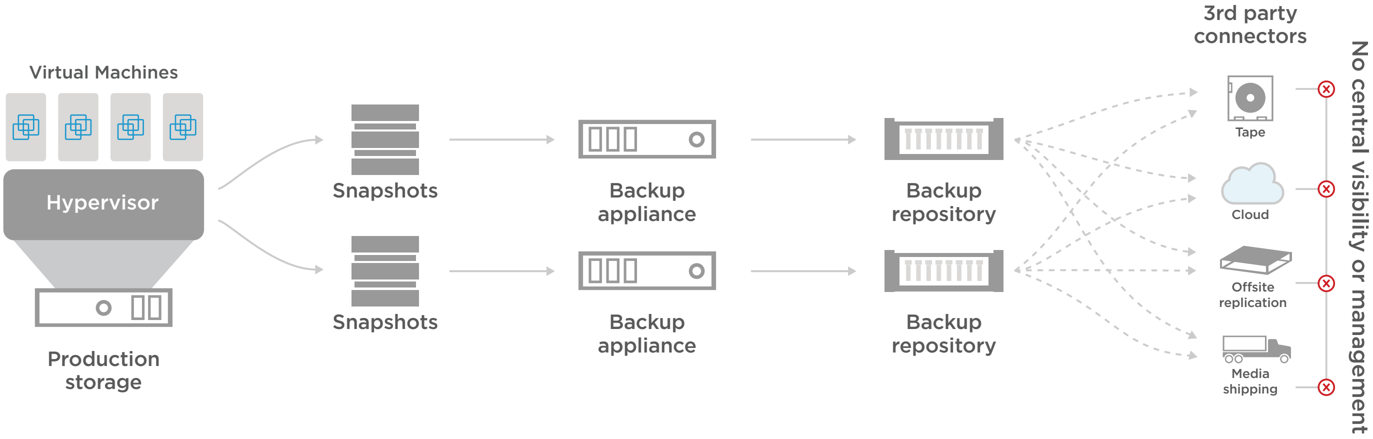

Different Companies, Different Needs

Every organization has its own unique requirements for the retention of virtualized data. For some, the requirements are defined by industry-specific laws and regulations; for others, the requirements are contingent upon the organization’s tolerance for risk and the potential damage caused by data loss. Regardless of their specific requirements, every company should have a set of clear, well-articulated policies and procedures regarding data protection, as well as minimum acceptable recovery point objective (RPO) and recovery time objective (RTO) times. Many organizations are looking for scalable SaaS solutions to address their needs.

Factors Driving the Need for a Single SaaS Data-Protection Solution

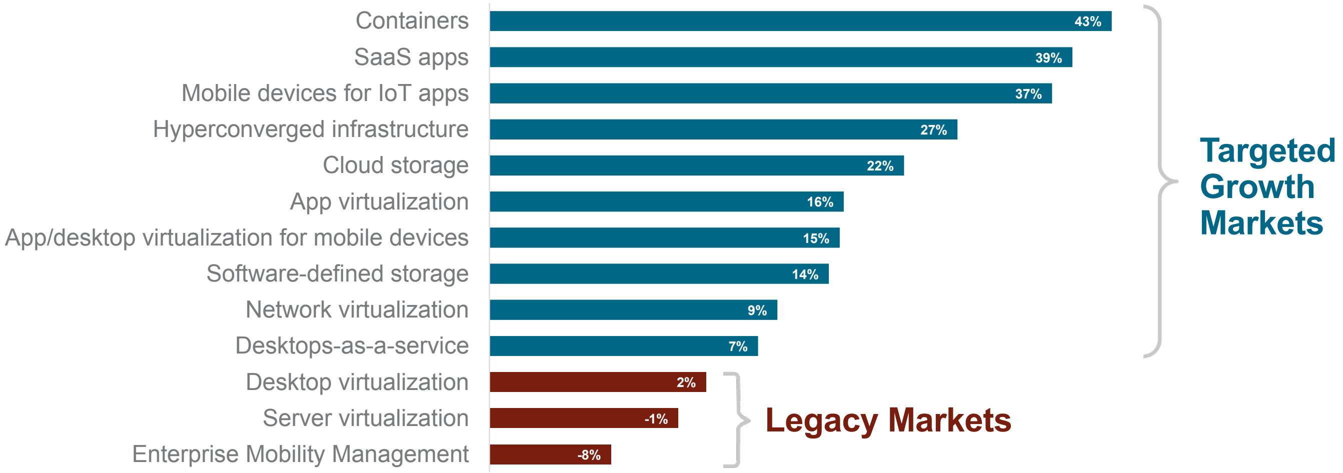

Source: TechTarget, 2017 Market Evolution and Repositioning Report

Recovering a virtualized environment from a legacy backup is complex, outdated, and expensive. The necessary media must first be recalled, potentially from a secondary remote site, and the data must be painstakingly restored to its original location. In the case of a disaster recovery (DR) scenario, which commonly occurs after a complete server failure, downtime may extend to days or even weeks with these types of legacy tools and processes.

Has Your IT Team "Gone Rogue"?

Traditionally, businesses have cobbled together multiple software solutions to address backup, archival, and DR as part of a larger data protection practice. In some cases, application, virtualization, and storage IT teams have become dissatisfied with traditional backup methods and now use the native point products available to them instead. They have, in effect, "gone rogue." In this environment, an organization's data-protection strategy often consists of a half-dozen or more different backup and recovery "solutions," which may or (most likely) may not be compatible with each other. The rise of this rogue IT fragmentation has resulted in protection silos and an "accidental" backup architecture (because no one would intentionally plan for an assortment of unconnected protection tools with no central oversight and no cost controls!). These accidental architectures increase cost, complexity, and risk as data protection environments scale.

Most Backup Solutions Are Not DR Solutions

It's important to note that most backup solutions are *not* DR solutions. In fact, these two solutions have traditionally demanded very different approaches. As a result, products that claim to manage both backup and DR have invariably proven to be unsuccessful at one or the other—or both. For example, backups are normally scheduled to occur at night (but not necessarily *every* night) to reduce CPU load and the associated negative impact on system performance. In this scenario, businesses need to be comfortable with a disaster recovery RPO approaching 24 hours. For many mission-critical applications, this is simply unacceptable from a business-continuity perspective.

Additionally, legacy products have traditionally employed tape-based systems, whether on-prem or remote, for data storage. This may be minimally viable for long-term archival storage, but it results in dramatically longer restore times and unacceptably long RTOs. Besides the restore speeds that tape typically offers, there is the additional time of retrieving tapes from an offsite vendor.

Organizations must also give careful consideration to the source of the data they are tasked with safeguarding. The use of SaaS applications in business has exploded in recent years, with 63 percent of companies now saying that more than one-fifth of the applications they use are delivered via the SaaS model—which is nearly double the number of SaaS applications in use five years ago. In such an environment, it is practical and preferable to backup SaaS data in the cloud. This allows companies to take full advantage of the speed, reliability, and geographic availability of public cloud providers such as Amazon Web Services (AWS).

Pervasiveness of SaaS: 2013 vs. 2018

Of all the applications used by your organization, approximately what percentage is currently delivered via the SaaS model? (Percent of respondents)

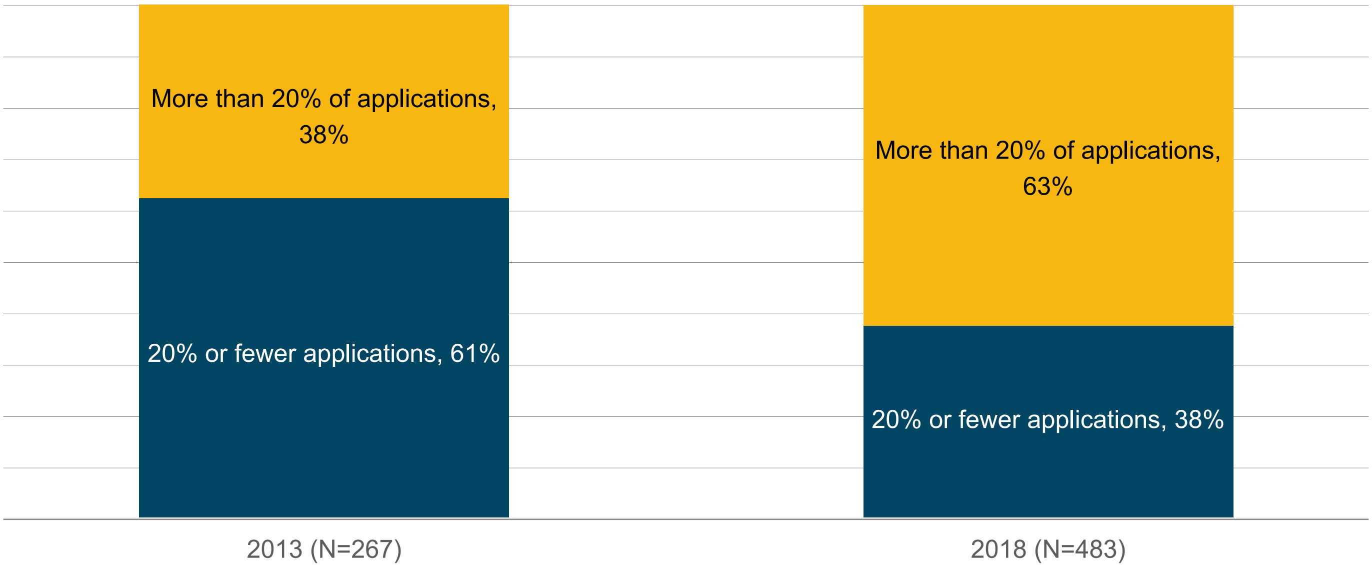

Source: ESG, 2018 IT Spending Intentions Survey

Given the inherent limitations of on-prem backup, companies who have, until now, trusted one legacy product to handle both backup and recovery duties have been gambling with their vital data, placing themselves at significant risk of lengthy downtime and potential loss of their data, revenue, and reputation. It's no surprise that these responsibilities—and ultimately, the blame for any failures—falls squarely on the shoulders of those directly in charge of maintaining and monitoring these solutions.

The Importance of Workload Mobility

It is becoming ever more common for companies to conduct business on a global scale. These organizations often find it necessary to replicate virtual appliances to geographically dispersed data centers for various reasons, including the following:

- Optimizing server speed and availability by situating workloads closer to clients
- Compliance with regional regulatory requirements
- Providing for redundancy and failover in the event of disaster
- Load balancing across multiple regions
- Safely testing and validating workloads without interfering with production environments

In each of these cases, it is critical that each data center is using the appropriate version of any virtual appliance and that any update or patch is replicated to every location that utilizes that application. Using the traditional model, physical media would need to be shipped across the globe, which carries the inherent risk of data becoming lost, damaged, or destroyed.

Cloud-Native Backup and Recovery for Virtual Environments

Today organizations are looking to the cloud for a viable, cost-effective alternative. But the true value of the cloud is often diminished by retrofitting legacy products into the cloud. This treats the cloud more like a tape drive or a hosting provider, rather than an efficient system unto itself. If you're going to use the cloud in their backup and recovery environment, you should use a system designed for the cloud and all its capabilities.

Cloud-Native SaaS Solution:

Backup Storage Initiatives Customers Plan to Deploy During 2017

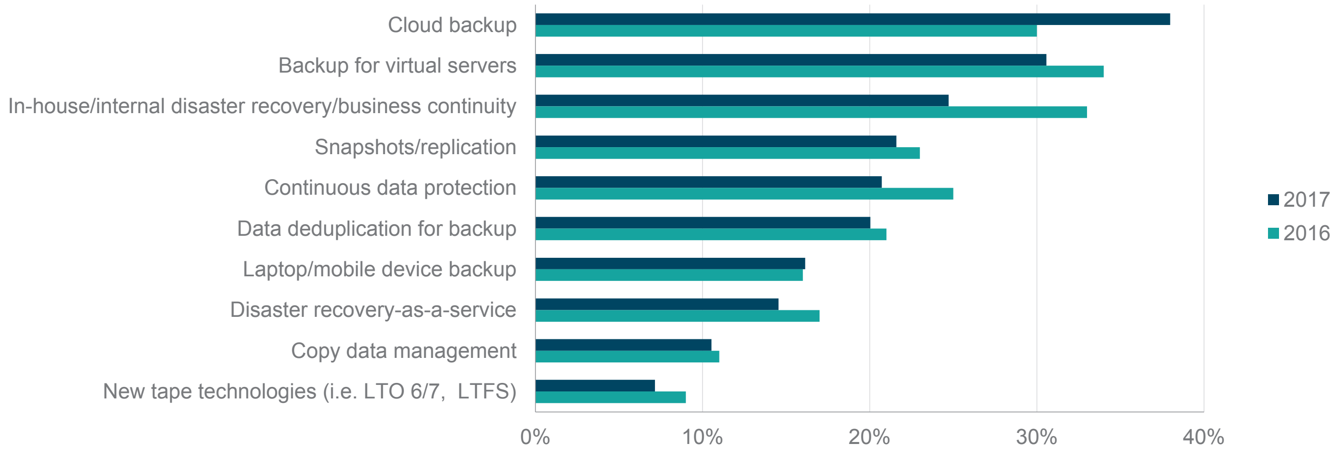

Source: TechTarget, 2017 Market Evolution and Repositioning Report

The threat of rogue IT, coupled with expensive hardware and the limitations of tape and disk-based storage means that it is no longer practical, or even affordable, for companies to host a complete backup, archival, and DR solution on-prem. But there's good news—with the maturation of the public cloud, there are now very real alternatives to on-prem backup, restore, and DR.

Here are five real-world reasons why you should consider moving your company's virtual environment backup, archival, and DR functions to a single, cloud-native platform:

1. Offsite Infrastructure

Cloud-based backup and recovery infrastructure is hosted completely offsite—there is no expensive hardware to requisition, and no software to license and keep up-to-date. Reliable, durable, fast, and cost-effective data recovery can therefore be enabled by an enterprise-grade, multi-region, public cloud infrastructure. This means your data is automatically protected without the need for administrative overhead. With virtual machines (VMs) replicated offsite, system downtime (and its resulting impact on productivity) can be reduced to mere minutes. Cloud-based systems also enable enterprises to store replicated VMs to multiple storage regions for even greater protection and redundancy.

2. Improved Business Agility

A cloud-native data management approach enables fast response times in the case of failover for DR, with RTOs measured in minutes. You can also move VMs across regions for regulatory needs or for development and/or test (dev/test) purposes. Businesses can back up their VMs from the data center to the public cloud, and move or recover entire VMs or individual files to any other region where they do business. Cost savings and speed and efficiency improvements are achieved by minimizing the storage footprint through global deduplication of data and optimizing storage tiers—with no extra infrastructure.

3. Dynamic DR and Workload Mobility

In the cloud, VMs can be configured for DR failover and be copied as needed for production or dev/test. Leveraging the public cloud also enables customers to be more flexible with their data movement. Workloads can be spread across geographies for easy replication and at-the-ready disaster recovery. Cloud-based DR eliminates the need for organizations to store complete copies of production systems at secondary company-managed data centers. Copies of VMs can also be pushed to any global location, making disaster recovery and workload mobility simple and efficient.

4. Simplified Management

Managing a cloud-based backup strategy is considerably easier than with on-prem systems. For example:

- Server backup and DR policies can be coordinated and monitored globally, from anywhere in the world, removing the burden of complex storage, compute, or networking management.
- Data tiering ensures that VMs are always stored cost-efficiently for long-term archival to address compliance demands—without the need for manual processes. Data is sorted into hot, warm, and cold tiers to optimize availability without adding unnecessary expense.
- Cloud-native content analysis capabilities provide IT administrators with a greater understanding of potential data and compliance risks across multiple data sources.

5. Radically Lower TCO

Unlike the capital expenditure (CapEx) model, the operating-expense (OpEx) model used by a cloud-based solution enables enterprises to pay only for what they use, thus improving decreasing the total cost of ownership (TCO). Not only does a cloud-based model eliminate costly hardware appliances and data centers and reduce administration needs, it can also provide a unified approach to backup, disaster recovery, archiving, and analytics on a single data set, which significantly lowers the costs created by data silos. In addition to auto-tiered storage, this approach provides a highly efficient collection of data in an ever-incremental backup model, avoiding the large volume of data stored by legacy models.

Additional Cloud Considerations

Governance and Compliance

Central auditability, legal admissibility, and long-term retention make compelling cases for easier data governance and compliance around data within virtual images. Since data is stored in the cloud, it is readily available for data mining, legal, and compliance needs. Enterprises can analyze backed-up data to understand the risks and challenges around dormant data, storage growth, and data classification. A cloud storage model can therefore increase visibility into existing data which can then be better leveraged for additional business value.

Development and Testing

Cloud storage allows for test/dev systems to be spun-up as needed, on demand, with no dedicated hardware or software. The result is greater flexibility and speed. By leveraging a copy of a VM in the cloud, IT teams can run tests and validation against a copy of the production data without disrupting critical production environments. A single protected VM can be centrally managed and replicated as often as needed, so development and testing can be easily managed around the clock and across geographies. VMs can also be repurposed at any time, so there's no need for separate systems.

The Big Takeaway

IT Investment Justification

Which of the following considerations do you believe will be most important in justifying IT investments to your organization's business management team over the next 12 months? (Percent of respondents, N=651, three responses accepted)

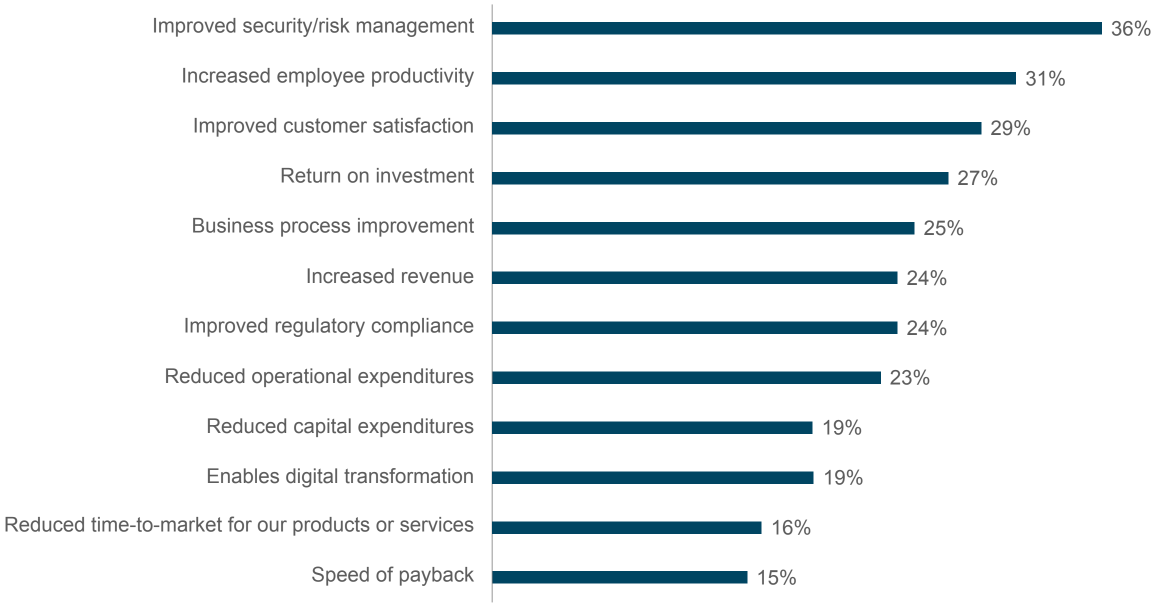

Source: ESG, 2018 IT Spending Intentions Survey

Many IT teams today are burdened with the unenviable task of trying to do more with less. The rise of virtualization and the increasingly distributed nature of business have resulted in an explosion of data, while budgets have stayed mostly flat—or increased only very slightly. Not surprisingly, making the business case for investment in IT infrastructure is often a difficult choice between competing priorities. Despite this, the security of a company's data must always be paramount, especially considering the stiff penalties introduced in Europe's General Data Protection Regulation (GDPR). As a result, 36 percent of companies say that improved security management is one of their top three spending priorities over the next 12 months.

Managing the backup and restore of VMs in a distributed environment is typically a convoluted process that involves multiple staff members supporting a complex architecture. Reduction in any new hardware, software, or administrative burden, therefore, improves business agility and radically lowers overall TCO.

To learn how Druva can help you achieve cost and time-savings by managing your VMware, VMware Cloud on AWS, Hyper-V, and Nutanix environments from within a single, centralized console, visit our virtualization solutions page.

About Druva

Druva is the global leader in Cloud Data Protection and Management, delivering the industry's first data management-as-a-service solution that aggregates data from endpoints, servers and cloud applications and leverages the public cloud to offer a single pane of glass to enable data protection, governance and intelligence—dramatically increasing the availability and visibility of business critical information, while reducing the risk, cost and complexity of managing and protecting it.

Druva's award-winning solutions intelligently collect data, and unify backup, disaster recovery, archival and governance capabilities onto a single, optimized data set. As the industry's fastest growing data protection provider, Druva is trusted by over 4,000 global organizations, and protects over 100 petabytes of data. Learn more at www.druva.com and join the conversation at twitter.com/druvainc.

Druva, Inc.
Americas: +1 888-248-4976
Europe: +44 (0) 203-7509440
India: +91 (0) 20 6726-3300
Japan: +81-3-6890-8667
Singapore: +65 3158-4985
Australia: +61 1300-312-729
sales@druva.com
www.druva.com